



An Audit of Facebook's Political Ad Policy Enforcement

Victor Le Pochat, *imec-DistriNet, KU Leuven*; Laura Edelson, *New York University*;
Tom Van Goethem and Wouter Joosen, *imec-DistriNet, KU Leuven*; Damon McCoy
and Tobias Lauinger, *New York University*

<https://www.usenix.org/conference/usenixsecurity22/presentation/lepochat>

**This paper is included in the Proceedings of the
31st USENIX Security Symposium.**

August 10–12, 2022 • Boston, MA, USA

978-1-939133-31-1

**Open access to the Proceedings of the
31st USENIX Security Symposium is
sponsored by USENIX.**

An Audit of Facebook’s Political Ad Policy Enforcement

Victor Le Pochat
imec-DistriNet, KU Leuven

Laura Edelson
New York University

Tom Van Goethem
imec-DistriNet, KU Leuven

Wouter Joosen
imec-DistriNet, KU Leuven

Damon McCoy
New York University

Tobias Lauinger
New York University

Abstract

Major technology companies strive to protect the integrity of political advertising on their platforms by implementing and enforcing self-regulatory policies that impose transparency requirements on political ads. In this paper, we quantify whether Facebook’s current enforcement correctly identifies political ads and ensures compliance by advertisers. In a comprehensive, large-scale analysis of 4.2 million political and 29.6 million non-political ads from 215,030 advertisers, we identify ads correctly detected as political (*true positives*), ads incorrectly detected (*false positives*), and ads missed by detection (*false negatives*). Facebook’s current enforcement appears imprecise: 61% more ads are missed than are detected worldwide, and 55% of U.S. detected ads are in fact non-political. Detection performance is uneven across countries, with some having up to 53 times higher false negative rates among clearly political pages than in the U.S. Moreover, enforcement appears inadequate for preventing systematic violations of political advertising policies: for example, advertisers were able to continue running political ads without disclosing them while they were temporarily prohibited in the U.S. We attribute these flaws to five gaps in Facebook’s current enforcement and transparency implementation, and close with recommendations to improve the security of the online political ad ecosystem.

1 Introduction

Online political advertising is a powerful tool for enabling engagement in the political process, but with this power comes the risk of abuse that can harm the integrity of the democratic process. Scrutiny of major online advertising platforms intensified due to foreign interference in the 2016 U.S. elections [87] as well as broader concerns on disinformation, voter suppression, and inauthentic behavior [89]. As government

Supplementary materials as well as an extended version of this paper that includes the appendices are available at <https://osf.io/7tw3e/>.

Table 1: Summary of (in)correctly classified ad counts, across undeclared political ads as labeled by Facebook or by us, within a 14-day observation period after an ad’s first activity.

Detected as political by Facebook		Not detected by Facebook
40,191 * False positive (subsection 5.2)	32,487 * True positive (subsection 5.1)	116,963 § False negative (subsection 6.3)
Not political		Actually political
<i>Precision: 0.45</i>		<i>Recall: 0.22</i> <i>F₁ score: 0.29</i>

* Across all advertisers worldwide; estimate based on 55% FP rate in U.S.

§ Across political advertisers worldwide.

regulation has failed to adapt [64, 71], oversight on online political advertising has fallen largely to the platforms themselves [42, 64]. Platforms therefore developed self-regulatory policies [71] that include verifying and revealing advertisers’ identity [48, 57], creating public archives of political ads [6], or even banning political ads altogether [1, 64].

A baseline requirement for platforms to protect integrity and reduce harm is then to properly identify advertisements that seek to influence public opinion, and adequately enforce their policies on those ads and their advertisers. Failing to do so correctly, rapidly, and consistently leaves an opportunity for ill-spirited advertisers to impede public scrutiny, spread violating content, and evade restrictions on political ads. Conversely, well-meaning advertisers are disadvantaged if their ads are unduly made unavailable due to incorrect enforcement, or if they (over-)comply with policies while others do not [72], especially when policies are unclear or ambiguous. Given the large number of submitted ads, platforms usually deploy automated methods for policy review, complemented by human review when needed [23, 38, 75].

In this paper, we audit whether Facebook makes accurate enforcement decisions for ads that may be in scope of its political ad¹ policies, but were not declared as such by the respec-

¹In this paper, we use ‘political ads’ as shorthand for ads in scope of Facebook’s policy, i.e., “ads about social issues, elections or politics” [11].

tive advertisers. Facebook is the most popular social media platform worldwide among users [66] and advertisers [103], and its political ad policies and transparency are considered to be among the most developed for major technology companies [71, 101], allowing us to analyze the effectiveness of self-regulation through one of the most advanced deployments. We build a novel large-scale data collection pipeline that retrieves all currently active ads running on Facebook’s core advertising platforms² from the *Ad Library*, its ad transparency tool. Our comprehensive and representative data set contains 4.2 million political and 29.6 million non-political ads from all 215,030 pages³ that ran political ads during the second half of 2020 and beginning of 2021, covering major elections in the U.S. and Brazil. We analyze the prevalence of ads that Facebook correctly detects to violate policies after they start running (*true positives*), ads that Facebook detects but are not political (*false positives*), and ads that Facebook fails to detect even though they are political (*false negatives*).

In prior work, the Ad Library has been used to study advertisers evading Facebook’s transparency requirements [36], while other research sought to quantify enforcement errors through anecdotal evidence [33, 97, 108, 111], or through crowdsourced [96, 102] or self-published [75] ads; however, these studies inherently cover only a small sample of ads. To the best of our knowledge, no previous study has quantified the *performance* of Facebook’s political ad policy *enforcement* in detecting non-compliance *at a large and representative scale*. A study such as ours is essential to understanding whether Facebook’s current self-regulation effort is sufficient to maintain the integrity of its political ad ecosystem.

Overall, we find that policy violations detected after an ad starts running represent a small share (1.7%) of political ads on Facebook. Detection happens rather quickly; yet it is worth noting that these violating ads failed to be detected during Facebook’s initial ad review, which allowed them to accumulate over 2 billion user impressions before being taken down. Unfortunately, this detection of violating ads seems to have little visible impact on advertisers. Despite a history of violations, we observe that the top violating advertisers were able to continue running new ads and accumulate more violations for long periods of time, even while political ads were banned in the U.S. [1, 70, 94].

Ambiguities in Facebook’s policies and flaws in Facebook’s existing detection appear to cause many unrelated ads to be incorrectly labeled as political: We estimate that among U.S. advertisers, 55% of ads detected as “political” by Facebook are in fact false positives. Conversely, we identify 39% of advertisers in Facebook’s Ad Library Reports as clearly political. While such advertisers are subject to a blanket rule in Facebook’s ad policy requiring them to declare *all* their ads as political, these pages ran a total of 116,963

ads that were not declared as political and not detected by Facebook. In addition, significant differences in the rates of undetected ads arise between countries: While performance is best in the United States at 0.85%, Facebook may fail to detect up to 45% of undeclared political ads in other countries.

When considering only the running ads where Facebook needed to make an enforcement decision, that is, ads not voluntarily disclosed by their advertisers, we find 61% more ads that are missed than are detected by Facebook within 14 days, and 55% of detected ads are likely false positives (Table 1). With more errors than correct decisions, Facebook’s current enforcement approach appears inadequate: users are left vulnerable to ads that seek to influence their opinion without proper disclosure, while legitimate advertisers regularly see their ads unjustly taken down. We attribute these flaws to insufficient attention for an advertiser’s political intent, the possibility for advertisers to continue running violating ads, inadequate localization in many countries, and ambiguity in policies, worsening transparency. Based on these observations, we make a number of recommendations to improve policy enforcement (section 7).

In summary, our main contributions are:

- We develop a novel data collection pipeline through which we obtain a comprehensive and representative view on all active political and non-political ads running between July 2020 and February 2021 across 215,030 pages (section 4).
- From an ad-level perspective, we find 1.7% of all “political” ads to have been detected post-hoc by Facebook, but detection is imprecise: we estimate through manual labeling that in the U.S., 55% of detected ads were incorrectly marked as political (*false positives*) and taken down (section 5).
- From an advertiser perspective, we find that detection of violating ads does not appear to prevent future violations, and that Facebook misses 116,963 ads from clearly political advertisers (*false negatives*), with considerably worse performance outside the United States (section 6).
- We identify five factors where our findings suggest that Facebook’s current enforcement and transparency implementation is lacking, and suggest improvements that would strengthen enforcement and improve the security of the online political ad ecosystem (section 7).

2 Background

2.1 Political ad policy

Facebook imposes increased authenticity and transparency requirements for “ads about social issues, elections or politics,” by requiring advertisers to confirm their identity and location and declare who funded the ads. These requirements are only mandatory and therefore proactively or reactively enforced in

²Facebook, Instagram, Messenger, and the Audience Network.

³An *advertiser* runs ads from their Facebook *page* [6]. In this paper, we use ‘advertiser’ and ‘page’ interchangeably.

around⁴ 60 countries and territories for ads about social issues, elections or politics, and in around 60 additional countries for ads about elections or politics only [18], with these sets of countries expanding over time. In all other countries, advertisers are currently “strongly encouraged” to get authorized and declare ads, but this is voluntary and not enforced [18].

Facebook considers ads to be “about social issues, elections or politics” if they are [12]:

- “made by, on behalf of or about a candidate for public office, a political figure, a political party, a political action committee or advocates for the outcome of an election to public office; or
- about any election, referendum or ballot initiative, including “get out the vote” or election information campaigns; or
- about any social issue in any place where the ad is being run; or
- regulated as political advertising.”

Facebook further specifies ‘social issues’ as “sensitive topics that are heavily debated, may influence the outcome of an election or result in/relate to existing or proposed legislation” and requires disclosure for these “social issue ads that seek to influence public opinion” [5]. Facebook defines a list of top-level ‘social issues’ per country (where applicable), which can change over time [5]; Facebook further clarifies these topics with examples of ads that are in and out of scope [56].

Before an advertiser can run ads about social issues, elections or politics in an applicable country, they must complete the *authorization process* there and confirm their identity and location [12, 24, 38, 48]. Once authorized, they can create ‘disclaimers’ to indicate which funding entity (individual, page or organization) paid for a given ad [22, 28, 38]. When running a political ad, the advertiser must then select it as a “Special Ad Category” [24] and add a disclaimer [24, 38].⁵

As shorthand, whenever we mention ‘political ads’ we refer to “ads about social issues, elections or politics” that were properly *declared* (i.e., having a disclaimer) or *detected* (i.e., lacking a disclaimer but marked as political by Facebook).

2.2 Policy enforcement

Facebook requires advertisers to self-determine that an ad is in scope of its ad policy on social issues, elections or politics, but also reviews any other submitted ad for policy compliance [38]. This “relies primarily on automated review (artificial intelligence) [...] and, in some cases, [they] have trained global teams to review specific ads” [38]. If an undeclared ad gets caught during this initial review, it never runs and is not archived in the Ad Library; the attempted violation will never be publicly known. This paper excludes such early detections.

⁴Lists of countries are inconsistent between the web portal [6] and documentation [18].

⁵Certain Facebook-vetted news publishers are exempt from declaring ads even if their content is political but not opinionated [7, 12].

If an undeclared ad passes review and is running, it can still “be flagged by AI or reported by [Facebook’s] community” as political [38]. Facebook then “disapproves” the ad retroactively, meaning they deactivate the ad, so it is no longer shown to any user. This is the type of ad detection that we study in this paper. Facebook also archives the violating ad in the Ad Library with a message that “this ad ran without a disclaimer,” regardless of whether the advertiser completed the ad authorization process [4, 6]. While the ad will remain publicly archived even when inactive, it will therefore never be known who paid for the ad. Violating pages may also be restricted from running new (political) ads or be disabled [12].

2.3 Transparency tools

Facebook emphasizes transparency as a means to hold them and their advertisers accountable [38], enabling users to be aware of who is trying to influence them as well as enabling journalists, organizations, and researchers (including us) to audit online political advertising [38]. To support this transparency, Facebook provides three core tools [113]:

1. The *Ad Library* [8] is a web portal where users can search all currently active ads for any Facebook page in any country, as well as all active and inactive ads about social issues, elections or politics. Only for the latter, provided metadata includes the disclaimer provided (if any), the identity of an authorized advertiser and how this was verified, and binned estimates of ad spend, reach, and impressions. A non-political ad disappears from the Ad Library once it becomes inactive; a political ad is archived for 7 years [6]. [Appendix A](#) shows how the web portal displays ads.
2. The *Ad Library API* [9] provides an interface for automated queries for all active and inactive ads about social issues, elections or politics for any page in a given country.
3. The *Ad Library Report* [10] aggregates advertiser data for all ads about social issues, elections or politics for countries where Facebook requires disclosure, listing all pages with at least one political ad in the chosen time span.

2.4 Related work

Prior work used crowdsourced or self-published ads to analyze the correctness of Facebook’s political ad policy enforcement. Silva et al. [96] developed a system to crowdsource Facebook ads in Brazil and classify them as political using a supervised machine learning model. Across 38,110 ads during the 2018 Brazilian elections, this model found 835 ads (2.2%) that had not been correctly declared nor detected as political. Matias et al. [75] conducted an audit study on Facebook and Google’s political ad policy enforcement through self-published ads, finding that Facebook applies their policies too restrictively, leading to 10 mistakenly prohibited ads (out of 238), while Google prohibited no ads. Sosnovik and Goga [102] compared platform, advertiser, and user perceptions of the def-

inition of online political ads on Facebook through 63,400 crowdsourced ads labeled by volunteers. They found that social issue ads in particular see the highest error rate due to unclear policies, although users largely perceive them as political, and observe disagreement between automated classifiers for political ads trained on differently sourced sets of (non-)political ads. Moreover, several media reports have given anecdotal evidence of ads missed by Facebook’s enforcement, both from politicians [33, 97, 108] and social issue organizations [97, 111]. Using the Ad Library, Cecere et al. [23] found that COVID-19-related ads were more likely to be detected by Facebook, suggesting that these may have been falsely detected, and that ad policies were confusing to advertisers. Our study quantifies the performance of Facebook’s enforcement on a larger and more representative scale than these previous studies, as we gather all *active* ads for *all* pages with at least one political ad, to analyze the prevalence of *both* false positives and false negatives.

From a transparency perspective, Edelson et al. [37] described and compared the efforts by Facebook, Google, and Twitter on a technical level. They later conducted a security analysis on transparency for Facebook’s Ad Library in the U.S. [36], finding that adversarial political advertisers could evade transparency requirements through erroneous disclaimers and undisclosed coordinated behavior. We assess how advertisers may evade declaring their ads as political altogether, through which they also avoid transparency.

Further audits of Facebook’s advertising platform found that advertisers can exploit ad targeting to infer private or sensitive user information [43, 50–52, 109], or to deploy highly targeted and biased ad campaigns [16, 49, 87], with users receiving inadequate targeting explanations from Facebook [16]. Facebook’s ad delivery may also skew which users see which ads, potentially leading to discrimination based on gender or race [14, 60, 69], including for political ads [15].

3 Enforcement Errors and Their Impact

We introduce the two error types that affect the security of Facebook’s political ad platform, i.e., the ‘threat model,’ and describe the actors that may either exploit these errors to induce harm, or that are themselves harmed by these errors.

First, ads may not be detected as political by Facebook, i.e., are *false negatives*. Once they are allowed to run, these missed ads harm the *integrity* of the online political ad ecosystem and of Facebook’s transparency efforts. They result from Facebook failing to discover ads that advertisers did not properly declare, whether deliberately to avoid scrutiny or accidentally due to misinterpreting (ambiguous) policies [33, 97, 102, 111]. Malicious advertisers may have an incentive to not declare politically motivated ads, as this relieves them of the accompanying restrictions. They would not need to get authorized by Facebook (requiring identification) nor display who paid for the ad [38]. Moreover, users will be unaware

that the advertiser is attempting to influence them as the ad interface will not reflect that the ad is political [38], and they might be shown the ad even if they requested to see fewer political ads [58]. The advertiser can then abuse these flaws to spread disinformation or prohibited content (e.g., voter suppression), or engage in ‘coordinated inauthentic behavior’ where accounts conspire to run influence campaigns [36, 89], without being publicly identified. Moreover, such an advertiser can circumvent bans on political ads, as was (temporarily) the case after the 2020 U.S. elections [1]. Finally, advertisers may want to evade transparency and accountability: undetected ads disappear from the Ad Library once they become inactive, leaving researchers and journalists unable to discover policy-violating content or hold advertisers and Facebook accountable for compliance with and enforcement of the political ad policy.

Second, ads may be incorrectly detected as political by Facebook, i.e., are *false positives*. As detected ads are taken down and may even result in pages being restricted from running ads or being deleted, Facebook reduces the *availability* of legitimate advertisements through these errors, whether the ads concern social issues (but do not influence public opinion) or are purely commercial. On the one hand, these can result from Facebook applying their policies too restrictively or erroneously and detecting ads that are in fact not political. For example, Facebook’s enforcement errors were found to hinder public health messages related to COVID-19 [23, 95] and vaccines [63, 79, 104]; were thought to unduly insinuate political division for social themes [47, 53, 80, 90, 99, 100, 114]; or resulted from false name matches [45, 74–76]. On the other hand, advertisers themselves may introduce false positives by (voluntarily) over-declaring ads that are not in scope of the political ad policy [102], possibly due to incorrectly or overly cautiously interpreting this policy or because they fear the ad will otherwise be erroneously detected and taken down. False detections also erode trust in enforcement as a whole, as they suggest that the automated decision models are unable to properly distinguish political ads, and further reduce the quality of input data to these models.

Throughout this paper, we label ads as follows:

Considered as political by			Label	Type
advertiser	Facebook	us		
✓	N/A	✓	declared	True positive
✓	N/A	✗	(over-)declared	False positive
✗	✓	✓	detected*	True positive
✗	✓	✗	(over-)detected*	False positive
✗	✗	✓	undetected*	False negative
✗	✗	✗	–	True negative
✗	(✓ OR ✓)		*undeclared	

4 Data Collection

To understand the dynamics and possible shortcomings of Facebook’s ad policy enforcement, we must capture the full lifespan for all relevant ads. The Ad Library API is insufficient for this purpose: it only returns ads *once* they are known to be political, which crucially excludes the period *before* Facebook enforces upon an ad, and omits ads that Facebook *never* enforces upon altogether. We therefore develop a novel large-scale data collection pipeline using the Ad Library web portal, which lists all active ads for a given page, regardless of whether they are political. In this section, we first define the scope of our study and present our data collection method. We then describe and validate the resulting data set, and discuss the ethics of our data collection as well as the impact of its limitations on our study.

4.1 Scope and method

Our data collection started on July 9, 2020, and was initialized with all pages that were present in any Ad Library Report for the past week since April 21, 2020, i.e., all pages that had published any political ads relatively recently. Until January 12, 2021, we continuously added pages newly appearing in the most recently available Ad Library Report. We consider the resulting set of pages as the scope of our study. We continued collecting ads for these pages for four more weeks, i.e., until February 9, 2021. This period therefore covered the pause on political ads in the United States after the elections on November 3, 2020 [1]. Our data collection covered all 71 countries with an Ad Library report at the time of our measurement. [Appendix G](#) lists the dates when reports were first available and when we started tracking their pages.

For every page that is in scope, we scrape its ads from the Ad Library web portal [8] with a target period of 24 hours, as well as page metadata with a target period of 14 days. We request all currently active ads which had impressions in the previous 7 days in any country; we do not apply any other filter. Additionally, for every ad, we gather its contents and metadata 14 days after its first observation through the ad snapshot tool used in the Ad Library API. As this endpoint reports the ad’s most recent state, even if already inactive, this allows us to observe any enforcement by Facebook within 14 days of the ad’s publication. For an ad detected within 14 days, we assume that ad detection led to its deactivation. We then calculate the activity period of an ad as the time between its first and last (daily) observation, assuming that the ad was published just before the former and detected just after the latter, with a 24-hour margin due to our scraping frequency. Moreover, whenever we analyze the activity period of an ad, we require that we likely captured the full lifespan of an ad, and therefore exclude ads active during the final four weeks of our data collection (reducing right censoring) or before/during our first scrape for a given page (reducing left

Table 2: Page/ad counts for the top 10 and other countries. For spend and impressions, we calculate the lower and upper bound based on the ranges available in our data.

Country	# pages ∇	with ad	# ads	for political ads		
				# ads	Spend (10 ⁶ USD)	Impressions (10 ⁹)
U.S.	90,018	69,815	21,934,716	1,902,473	810–1,146	45–53
Brazil	39,675	33,459	1,039,109	696,612	13–27	8.4–10
India	13,798	10,722	779,670	121,269	2.9–3.6	3.9–4.4
Italy	11,758	10,049	384,808	124,767	6.5–22	2.7–3.2
U.K.	10,558	8,016	2,456,921	96,660	17–33	2.4–2.8
Germany	10,223	8,501	782,078	121,482	13–30	2.8–3.2
Ukraine	8,632	7,437	315,425	133,006	3.2–17	3.4–4.0
Mexico	7,229	6,145	275,028	88,406	4.9–6.9	3.8–4.4
Canada	6,352	5,198	529,107	76,076	13–22	1.8–2.1
Romania	5,708	4,927	189,389	104,228	5.8–9.1	2.8–3.3
Other	61,873	50,761	5,142,518	726,382	54–109	18–21
Total	265,824	215,030	33,828,769	4,191,361	944–1,426	95–112

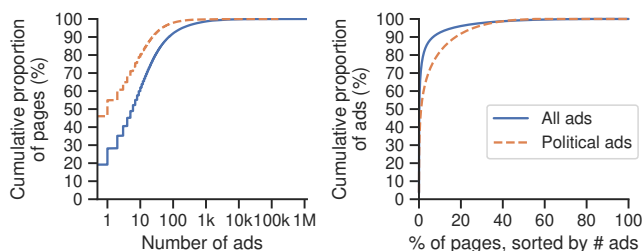


Figure 1: Distributions of observed ads and pages.

censoring). Finally, we retrieve all political ads that were active during our measurement period through the Ad Library API on March 30, 2021, in the subset of countries where API data was available to us (covering 80% of scraped ads, [Appendix G](#)).

4.2 Data set description

In total, we observed 33.8 million unique ads during our measurement, of which 4.2 million were declared or detected as political (Table 2). As Facebook only provides ranges for ad spend and impressions for political ads, we estimate that these had around 100 billion impressions and cost around 1 to 1.4 billion U.S. dollars. We see that United States advertisers dominate our data in terms of the number of ads placed overall, as well as in political ad count, spend, and impressions until the 2020 U.S. elections ([Appendix B](#)), when Facebook restricted U.S. political ads [1].

We observed ads across 215,030 pages (by definition, all of these pages had at least one political ad ever); we never observed any ads for 50,794 additional pages in scope. Small advertisers represent the majority of these pages: the median page posts fewer than 6 ads and 1 political ad respectively ([Figure 1](#)). Conversely, a small percentage of pages account for the majority of ads: the top 20% advertisers posted 92.5% of all ads and 81.9% of all political ads. Distributions are similar between U.S. and non-U.S. advertisers. To conserve

resources, we manually analyzed the largest advertisers up to then on four occasions (Appendix B), and discarded those that we considered unlikely to intentionally publish political ads. Before discarding, these advertisers placed 5,862,808 ads (17.3% of all ads). Our analysis in Appendix C confirms that these pages minimally placed political ads (0.011% of their observed ads), supporting our decision to discard them.

4.3 Data set validation

We assess the coverage of our data set both internally and externally to determine how often we were unable to retrieve all available ads. As an internal validation, we compare the expected number of available ads included in Facebook portal data to the number of actually observed ads. As an external validation, we compare the observed political ads with those retrieved from the Ad Library API.

We summarize our coverage in Figure 2. We missed 19.8% of ad observations, most often due to a limitation in Facebook’s systems: even though the portal is able to report that a page has over 50,000 ads, it fails to actually load more than 7,800 ads per scrape. Very large advertisers therefore bear the bulk of missing observations. We also miss the first ads from newly seen advertisers due to a delay of usually three days between a page’s first political ad impression and its appearance in the Ad Library Report [32]. Otherwise, discrepancies are due to our scraping frequency or setup: we miss ads that disappear during a scrape, that only appear between scrapes or after the last scrape, or when the scraper (partially) failed. For 7.5% of ads, we could not make the 14-day snapshot: this is largely due to resources being unavailable or restricted through Facebook’s snapshot tool, or because the 14-day interval was outside our measurement period. Finally, for 184 pages (0.07%), we failed to retrieve page metadata.

Based on the maximum number of missed observations per page, we estimate to have missed at least an additional 6.4% of ads. Weighted by the observed proportion of political ads per page, we estimate to have missed 11.4% of political ad observations. We note that the (unknown) number of unique ads that we were unable to retrieve is significantly lower than the number of missed ad observations, since many ads are active for more than one day. Finally, based on the API data, we missed around 1 million political ads (24.8%), with an estimated combined spend between 188 and 550 million USD, and 24 to 31 billion impressions. While missing data may introduce risks to research validity [46], our findings are lower bounds mainly calculated in the aggregate, which are less affected by our data gaps. We therefore believe that our data and results remain sufficiently representative for the Facebook political ad ecosystem.

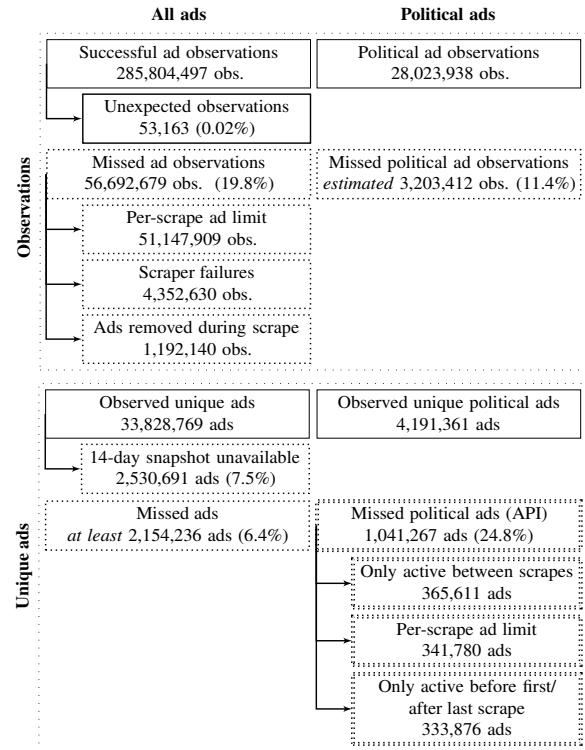


Figure 2: Summary of data set coverage.

4.4 Ethics

We follow ethical guidelines for Internet measurement and cybersecurity research [19, 54, 73]. Our data collection does not affect any non-advertiser Facebook users and we do not observe any personally identifiable information on them. Our research received an IRB exemption as it does not involve human subjects. As part of publicly available metadata, we collect the disclaimer that the advertiser provides to Facebook for the ad authorization process [48]; in the case of individuals, this may include personally identifiable information such as their legal name. We only process this data in the aggregate and do not use it to identify any individual. Similarly, we do not name any specific advertiser to avoid inflicting harm resulting from flaws in Facebook’s enforcement.

We collect only publicly accessible data. As Facebook states that “more than 2 million people visit the Ad Library every month” [3], we do not expect this data collection to significantly affect the availability of the Ad Library, and we did not observe any service outage possibly caused by our scraping. While Facebook’s ‘Automated Data Collection Terms’ [17] may prohibit scraping, we believe that our research is in the public interest, and that its societal benefits justify the technical resources consumed from Facebook, as well as potential reputational and financial harm on Facebook. Institutions, civil society organizations, and researchers have previously called for improved transparency for all ads [2, 26,

35, 39, 59, 62, 64, 65, 71, 88, 110], confirming the value of our data set. We share our data and methods with other researchers at <https://osf.io/7tw3e/>. In the context of prior work, we already communicated with Facebook to discuss their ad review and transparency, and presented to them the overall issues and recommendations that we also analyze in this paper.

4.5 Limitations

The definition of our scope leads us to only track pages with at least one known political ad (declared or detected) as recorded in the Ad Library Report. If a page never gets caught or is exempt (news pages [7]), we will therefore not track its ads, potentially missing their false negatives. Likewise, we cannot cover advertisers in countries where declaration is voluntary, as Facebook does not publish an Ad Library Report there [6].

As quantified in subsection 4.3, we do not achieve full coverage of the ads published by advertisers within our scope. On the one hand, coverage is affected by the trade-off between limited resources on our side and the large number of pages in scope. The 24-hour scraping period means we miss ads that were active only between two scrapes, and limits the granularity of observed activity periods. Our snapshot delay leads us to miss ad status changes beyond 14 days. However, we consider any changes after 14 days less likely to be due to Facebook’s own enforcement system,⁶ but rather due to external reporting. We also request ads shown in any country; Facebook provides a filter by country, but this would prohibitively multiply the required resources.⁷ We therefore assign pages and their ads to a country based on a heuristic, i.e., top spend on political ads. On the other hand, delays, flaws and changes in Facebook’s systems further reduce coverage. Delays in the publication of the Ad Library Report [32] and a limit on the number of retrievable ads cause us to partially miss ads from newly added and very large advertisers, respectively. We also experience infrequent failures of our scrapers, due to changed request methods or unavailable resources, or race conditions during one scrape (e.g., leading to duplicate ads). Beyond these unobserved ads in scope, we do not know the total number of ads on Facebook, which prevents us from quantifying true negatives and calculating metrics that depend on it. However, we expect true negatives to be much more prevalent, and therefore select classification metrics that are more robust against this class imbalance.

Finally, limitations result from Facebook’s transparency implementation. Without full metadata on all ads, we cannot quantify the impact in terms of spend and impressions of undetected political ads. We also have no visibility into ads that are caught during initial review and are therefore pre-

⁶Facebook started its pause on political ads 7 days before the 2020 U.S. election, a.o., to “re-review[ads] for policy violations” [107].

⁷However, we find that this filter can also be unreliable, with some ads only being available when no country filter is set.

vented from running altogether. More abstractly, we rely on Facebook’s Ad Library functioning properly, i.e., returning the actual, complete set of (non-)political ads from all pages [21, 110, 112]. While we have no reason to believe this is not the case, we also have no way to confirm this for our data set, due to a lack of transparency into Facebook’s architecture. Crowdsourced ads may allow to audit the accuracy of the Ad Library, albeit not completely [21, 30]. Moreover, while we are the first to conduct large-scale data collection through the web portal, researchers and organizations have documented consistency, completeness, and reliability issues with the Ad Library API [22, 25, 32, 36, 39, 41, 91, 93]. These challenges in comprehensively obtaining all currently active ads ultimately harm Facebook’s transparency efforts.

5 Ad-level Enforcement

We first examine enforcement of individual ads, independent from the advertiser. We start by quantifying the prevalence of enforcement, that is, how frequently Facebook is taking down ads for not having the required disclaimers, and determine the exposure that violating ads had before detection. We then survey how often an enforcement decision made by Facebook is appropriate, especially with regard to ads that likely should not have been taken down (*false positives*) and where incorrect enforcement harmed advertisers.

5.1 Current ad policy enforcement

Within our measurement data, 72,678 ads were marked at some point as ‘detected,’ i.e., political but not properly declared, within 14 days after the ad’s first activity. These detected ads therefore represent a minor share (1.7%) of all 4.2 million observed political ads. Edelson et al. [36] reported a 9.7% detection rate for May 2018–June 2019, suggesting policy awareness and compliance may have since improved. Moreover, Facebook has stated that “between March 1[, 2021] and Election Day, [they] rejected about 3.3 million ad submissions that targeted the US without completing the authorization process before they could run” [3], suggesting that Facebook’s initial ad review already catches most violations, although the lack of detail makes a reliable comparison difficult (e.g., authorization may be subject to separate review, and an advertiser could try and resubmit an ad until it passes).

Next, we analyze whether Facebook prevents violating ads from gaining much exposure by measuring how quickly Facebook takes down an ad that should have been declared as political. Detection of ads that slipped through the initial ad review is relatively fast (Figure 3): 40% of ads were detected within less than 1 day, with the median activity period being less than 2 days. Detected ads are also active for shorter periods than any political ad, for which the median activity period is less than 3 days, suggesting that ads are primarily detected while they are still active. However, violating ads may still

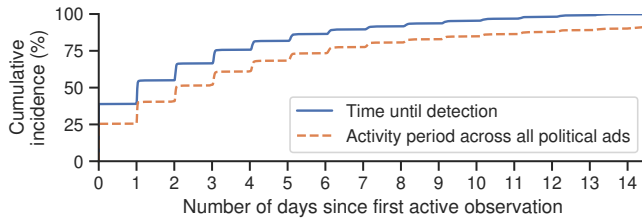


Figure 3: Activity period for detected and all political ads where we likely observed the full lifespan (subsection 4.1) and observed detections within 14 days.

enjoy significant exposure in budget, impressions, and time. We estimate the detected ads to have accumulated spending between 12.2 and 20.7 million U.S. dollars and between 2.1 and 2.4 billion impressions, i.e., instances where a user saw the ad without the proper context that it was political. 5,885 ads (8.1%) were active for over a week, meaning detection occurred very late. Moreover, we find 49,263 ads that were likely detected only after they became inactive, as they became inactive within 14 days, were not yet marked as political after 14 days, but were present in the Ad Library API. The advertiser could therefore display their violating ad for the desired duration. These 55,148 combined detections do not prevent most or any user harm, as most or all intended ad impressions still occur. Instead, they are only useful for the secondary goal of transparency (as the ads are then included in the Ad Library) and for any potential disciplinary measures taken against the page.

5.2 Ads incorrectly detected as political

When Facebook takes down ads for a lack of disclosure of their political nature, some of these decisions are incorrect, i.e., *false positives*. For example, the takedown of 1,413 ads (1.16% of all detected ads) was later undone, possibly after an appeal by the advertiser. These reflect admissions by Facebook that the ads were false positives and should not have been disabled. To study false positives more systematically, we labeled a randomly selected sample of 300 correctly declared and 300 detected ads. We restricted these samples to advertisers in the United States to ensure that annotators could interpret ad text and context. Three authors determined whether each ad was within or outside the scope of Facebook’s political ad policy. They were instructed to adhere as closely as possible to Facebook’s definition, i.e., not apply their own interpretation of what should be a political ad. In a subsequent meeting, the annotators revisited disagreed-upon ads, and reassigned a final agreed-upon label in the case of simple labeling errors (agreement on the definition, but for example a missed reference to a politician). Otherwise, if they considered Facebook’s definition too ambiguous, in particular on whether an ad sought to influence public opinion, they recorded “disagreement” as the final outcome. Using Krip-

endorff’s α [68], we achieve an inter-rater reliability of 0.81, i.e., sufficiently strong agreement for reliable conclusions.

Table 3 lists the results of our labeling. For *declared* ads, a false positive indicates that an advertiser unnecessarily declared that ad. Across our sample, we observe 3.3% over-declared ads, suggesting the practice is rare. 80% of declared ads are related to politics and elections, which are clearly in scope of Facebook’s ad policy and should therefore be declared. Across all observed ads, Facebook does not appear to retroactively mark declared ads as non-political, i.e., Facebook does not (need to) check whether a declared ad falls within the scope of its policy. For *detected* ads, a false positive indicates over-enforcement by Facebook, unduly taking down the ad. Across our sample, a majority of detected ads (55%) should not have been enforced upon; if we extrapolate this rate to all detected ads, 67,433 ads should not have been taken down. This suggests that Facebook’s ad detection is overly aggressive. Edelson et al. [36] observed a 79% false positive rate through a similar manual analysis, corroborating our finding that this rate may be very high.

The annotators also described the ad topic, using an inductively developed codebook that was aligned in their meeting (Appendix E). 24% of false positive detected ads concern commercial products or services, and it was not immediately obvious why Facebook detected those ads as political. Among ads where the likely cause of error was more discernible, most errors concerned COVID-19-related and health ads, for which the ambiguity in the definition of ‘social issues’ (“seek to influence public opinion”) makes confusion for advertisers and Facebook’s review more understandable. It appears that these policy ambiguities account for most errors; we attribute only 6 false positives to a likely false “keyword” match (e.g., a shop called ‘Mayors’). Matias et al. [75] observed a 4.2% false positive rate across ambiguous issue ads and false matches. Our findings suggest that in addition, commercial ads, where the reason for detection is less clear, represent a significant share of false positives. Finally, the annotators could not agree on the label for 5% of declared and 23% of detected ads, highlighting the difficulty of interpreting Facebook’s ad policy consistently. This disagreement mainly involved ads relating to social issues, where it was unclear whether the ad sought to influence public opinion, or ads with incidental references to politics, such as a candidate’s yard sign being visible in a real estate listing.

Summary We find that detected political ads without a disclaimer account for only a small share (1.7%) of observed political ads, and that Facebook detected them rather quickly. Still, we see 55,148 detected ads running for more than a week or for their full intended duration, making detection largely ineffective at preventing users from seeing these violating political ads. In addition, detection appears to be very imprecise: we find that 55% of detected ads in the U.S. should not have been taken down (*false positives*), harming advertisers by making their legitimate ads unavailable to users.

Table 3: Manual categorization of 300 declared and 300 detected ads, grouped by annotators’ assessment of whether these are political per Facebook’s ad policy. ⊙ : Related to social issues. Percentages are given within the sets of declared and detected ads, respectively; the margin of error is for a 95% binomial proportion confidence interval.

Ads considered political (<i>true positives</i>)					Ads considered non-political (<i>false positives</i>)					Ads disagreed upon by labelers				
Topic	declared		detected		Topic	declared		detected		Topic	declared		detected	
	#	%	#	%		#	%	#	%		#	%	#	%
By a political figure/organization	143	47.7	1	0.33	Commercial product/service	0	0.00	73	24.3	Environment ⊙	4	1.33	13	4.33
About a political figure/organization	61	20.3	15	5.00	COVID-19-related	0	0.00	24	8.00	Insignificant reference to politics	2	0.67	9	3.00
About elections	35	11.7	13	4.33	Health ⊙	5	1.67	18	6.00	Food assistance ⊙	2	0.67	6	2.00
Political Values and Governance ⊙	10	3.33	10	3.33	News/media	1	0.33	8	2.67	News/media	0	0.00	8	2.67
Civil rights ⊙	5	1.67	13	4.33	Apps/games/websites	0	0.00	8	2.67	COVID-19-related	1	0.33	6	2.00
Environment ⊙	6	2.00	4	1.33	Arts/Personalities	0	0.00	7	2.33	Health ⊙	2	0.67	4	1.33
Economy ⊙	6	2.00	3	1.00	Other	0	0.00	24	8.00	Other	1	0.33	3	1.00
Other ⊙	9	3.00	6	2.00	Other ⊙	4	1.33	4	1.33	Other ⊙	3	1.00	20	6.67
Total (Precision)	275	91.7	65	21.7	Total (False discovery rate)	10	3.33	166	55.3	Total (Disagreement rate)	15	5.00	69	23.0
Margin of error		±3.1		±4.7	Margin of error		±2.0		±5.6	Margin of error		±2.5		±4.8

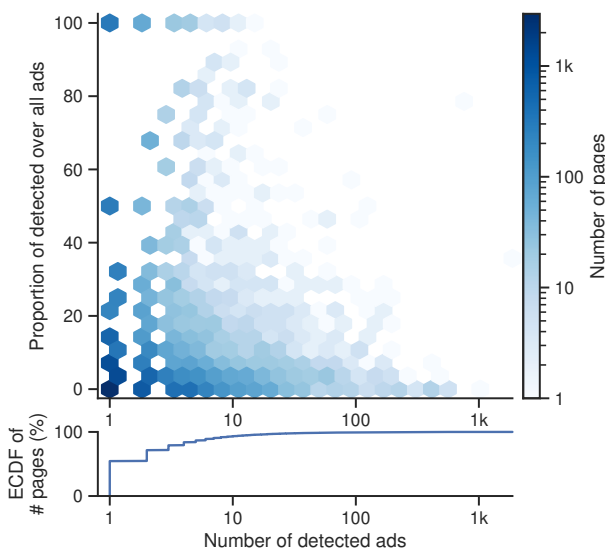


Figure 4: Detected ads versus their share of all ads for a page.

6 Page-level Enforcement

We continue at the level of a Facebook page to examine whether enforcement appears to be consistently and correctly applied across all ads of a page. We start by describing how advertisers react to takedowns of a page’s ads. We then classify pages to describe the current composition of Facebook’s Ad Library based on a page’s likely political intent, and identify likely *false negatives* as ads published by pages with a clear political purpose that Facebook failed to detect.

6.1 Reaction to enforcement

We first analyze whether advertisers are able to repeatedly violate Facebook’s ad policy. We observe detections of undeclared political ads for 13,900 pages (5.2%), again lower than the 68.3% rate observed by Edelson et al. [36] for May 2018–June 2019. Overall counts of detected ads per page

were low (Figure 4, left). However, 7,535 pages (54.2%) did not declare any political ads, so their only political ads were those that were detected (Figure 4, top), suggesting they may have unintentionally posted ads that were deemed to be political, possibly due to insufficient awareness of or ambiguity in ad policies. No advertisers appear to have placed many and mostly political ads without declaring them (Figure 4, top right); high absolute counts of detected ads are an artifact of an overall high volume of ads (Figure 4, bottom right). It does not appear that Facebook frequently banned pages as an enforcement action after detecting undeclared ads: only 458 pages were deleted some time after an ad detection, or 3.3% of pages with detected undeclared ads. (For reference, 7.2% of all pages in scope were ever deleted.) 373 of these 458 pages even continued placing new ads between their last taken down ad and deletion of the page.

Next, we analyze whether advertisers declare more ads after ads have been taken down, i.e., whether Facebook’s enforcement increased adherence to its ad policies. Ideally, this reaction should prevent future violations and protect users from being exposed to unmarked political ads. Out of the top 75 pages ranked by detected ad count (listed by class in Figure 5), 22 increased their proportion of ads declared as political after being detected (①–④): 5 started declaring continuously (①) and 5 others only shortly did not declare (②). However, increased declaration was only short-lived for 12 pages (③–④). Furthermore, 48 pages (⑤–⑥) barely declared any ad as political and often had a steady stream of violating ads taken down by Facebook. (This includes 39 news aggregator pages (⑥) that are likely not exempt from declaration, unlike more traditional news organizations [7].) This suggests that the most frequent offenders (in absolute terms) did not face any durable restriction in their ability to run ads as a potential disciplinary measure imposed by Facebook to increase compliance. Despite sometimes frequent and prolonged violations, all pages in the top 75 continued publishing ads after detection.

After the 2020 U.S. election, Facebook temporarily disallowed political ads on their platform [1]. Instead of ceasing to

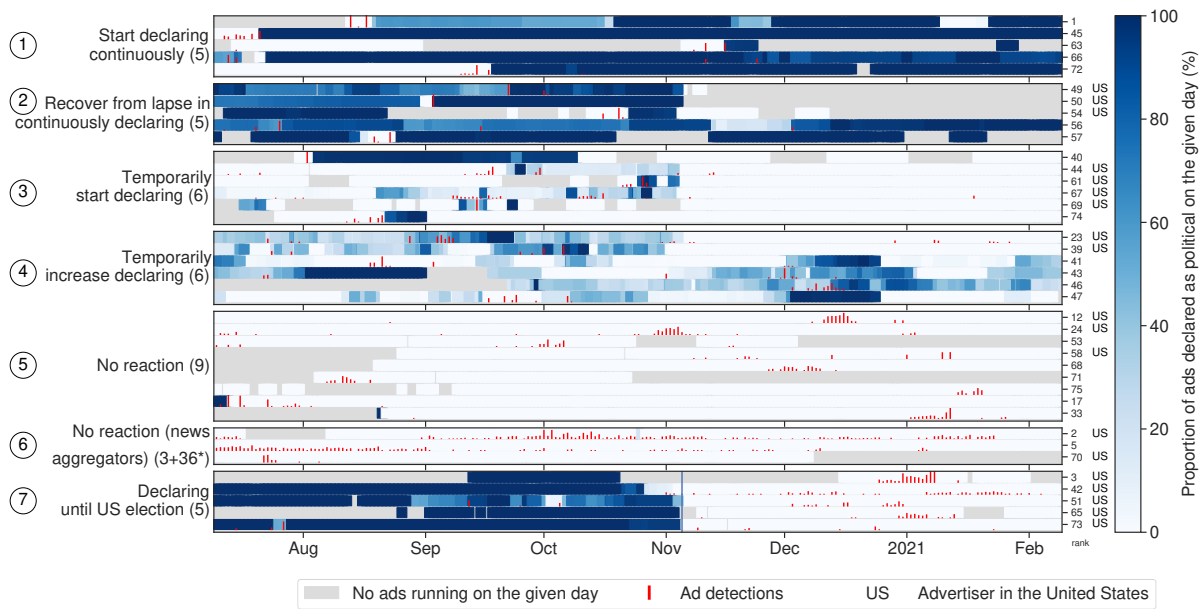


Figure 5: Reactions to ad detection by the 75 advertisers with most detected ads. Each row corresponds to an advertiser, with shades of blue indicating the proportion of ads the respective advertiser declared as political over time (aggregated daily). Red markers at the bottom of each row denote when ads were detected and taken down by Facebook. Advertisers are grouped into seven classes (denoted by ⑦) based on their reaction to enforcement (start/increase/resume declaring political ads) and the duration of increased declaration (continuously/temporarily/no reaction). For example, group ⑥ contains advertisers that have a steady stream of ads taken down by Facebook but do not show any apparent reaction to enforcement (they continue not declaring any ads as political). These advertisers are news aggregator pages; we show three examples and omit 36 similar pages (*).

advertise, 5 of the top 75 violating pages continued running ads but *stopped* declaring them as political (⑦), even though they were clearly of political nature as they had previously declared (nearly) all their ads as political. Even though these 5 pages advertised merchandise such as T-shirts with political messages, or were advocating for civil rights and environmental policy, Facebook only detected and took down 3% of their ads running after the election. Overall, 1,018 pages ran 71,426 undeclared ads after the U.S. election, whereas they only ran declared political ads before then. These pages did not appear to be deterred by the political ad pause, and Facebook did not effectively prevent them from running ads that were very likely political. This failure of enforcement rendered the ad pause less effective, and put these violating pages at an unfair advantage over advertisers who did comply and ceased running political ads as required [70, 94].

6.2 Current enforcement by advertiser class

Facebook’s policy on ads that require disclosure is broader than just ads published by obvious political actors such as parties or candidates [64]. Consequently, the Ad Library Report also lists advertisers beyond those actors, such as those placing ads about social issues, in partnership with (on behalf of) political actors, or with (non-partisan) calls to vote, next to advertisers with incorrectly detected ads (subsection 5.2).

6.2.1 Page classification

To quantify the prevalence of different types of advertisers, we match internal and external data sources with observed Facebook page metadata to classify pages into one of four topics. If a page is listed in multiple sources, we select its main class using the following order: (1) political, (2) government, (3) media, and (4) issue-related.

- We retrieve political committees for the 2020 U.S. elections from Facebook’s Ad Spending Tracker [106] (matched on page name) and the OpenSecrets project of the Center for Responsive Politics [85] (matched on page ID).
- We retrieve political candidates and parties registered for the 2020 Brazilian municipal elections from the Superior Electoral Court [98] (matched on CNPJ).
- We retrieve pages who identified themselves during Facebook’s advertiser authorization process through a Federal Election Commission identification number (FEC ID) for U.S. political pages or U.S. government credentials for U.S. government pages from our data (matched on page ID).
- We retrieve media organizations from Media Bias/Fact Check [77] and NewsGuard [83] (matched on page alias), with local pages for large news aggregators manually added.
- We retrieve U.S. nonprofit (tax-exempt) organizations as potential issue-related pages from National Center for Charitable Statistics [20, 81] and Internal Revenue Service [40]

- data (matched on disclaimer and ZIP code).
- We retrieve manually curated *Explore* lists [86] containing political, government, media, and issue-related pages from CrowdTangle [29] (matched on page ID).
- We enumerate the most common Facebook page categories for pages within the previous data sets, manually select those categories that are sufficiently specific to a class, and then retrieve all pages within those categories (Appendix F) from our data (matched on page ID).
- We retrieve all pages that completed Facebook’s advertiser authorization process from our data (matched on page ID).

Overall, 59.7% of pages fall into one of the four categories. We distinguish an additional 11.3% of pages outside these topics that completed the authorization process, as this implies a genuine intent to at least sometimes place political ads.

6.2.2 Distribution of political ads over classes

We first analyze the composition of advertisers listed in the Ad Library Report globally, i.e., within our scope of pages with at least one recent (declared or detected) political ad (subsection 4.1). We expect mostly political and issue pages to appear in the Report; however, based on our classification (Table 4), ‘obvious’ political advertisers only represent 39% of measured pages, with a further 8% that are issue advertisers; these combined account for 73% of observed political ads. Only 53% of all pages were authorized to declare ads as political. Unauthorized pages may have no political motive, suggesting they inadvertently published ads that fall under the political ad policy,⁸ or their ads were incorrectly detected as political by Facebook. Alternatively, they may be political actors that refuse to authorize themselves, or may be unable to do so due to Facebook’s policies, e.g., if they are outside the country in which they want to run political ads [36].

We further analyze whether certain page classes are more likely to have such “unintentionally” undeclared and detected ads by comparing *detected* with *overall* ad counts (Table 4). Government (6% detected vs. 4% overall) and issue (16% vs. 8%) pages are overrepresented, hinting at discrepancies between their and Facebook’s understanding of which ads should be declared. Media pages account for the most detected ads in absolute numbers (34%), but place ads in similarly high volumes (35% of all ads). 46% of authorized pages and 21% of political pages failed to declare at least one ad that was later detected, even though Facebook’s ad policy requires all ads from or on behalf of political figures to be disclosed.

Next, we measure the proportions of political ads over all ads per page and class. If an advertiser is political in nature, we expect them to have 100% political ads, either because they properly declare all their ads or because Facebook detects their undeclared ads. Indeed, this largely holds for identified political pages (Table 4), where 81% had only political

⁸These ads then likely triggered inclusion in the Report; due to delays in this inclusion, we cannot observe these one-off political ads (subsection 4.5).

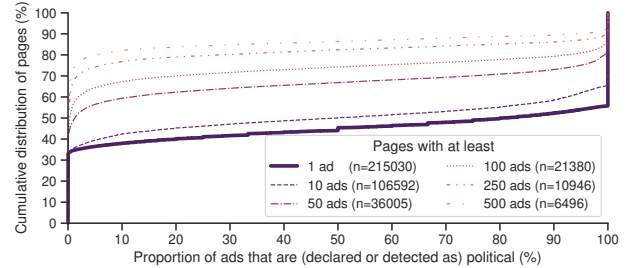


Figure 6: Distribution of political ad proportion over pages.

ads. However, for government, issue, and media pages, this share is much lower, at 23%, 22%, and 12% respectively, and median proportions of political ads of 20%, 25%, and 0% respectively, showing that the Report contains many (classes of) pages whose ads are mostly non-political. Across all pages with observed ads (Figure 6), we similarly see that 44% had only political ads, while 33% had no political ads observed during our measurement, with the latter increasing for larger advertisers, suggesting these may not have political intent.

Overall, we find that 47% of pages in the Ad Library Report are not authorized, with over 33% of pages hardly publishing political ads over time. These pages may have had incorrectly detected ads (*false positives*) or placed an ad that they did not consider political even though Facebook did. Indeed, government and issue pages are much more likely to have ads detected by Facebook. When pages have no clear political motive, the advertisers as well as Facebook must determine at the individual ad level whether the ad is in scope of the political ad policy, which may be more prone to interpretation errors and disagreements and therefore lead to enforcement errors. Conversely, we consider 39% of pages to be core political actors. Although Facebook’s policies require any ad made by a political actor to be declared [11], enforcement is still necessary as 21% of such pages have at least one detected ad. Next, we analyze whether these pages had any ads that were neither declared nor detected by Facebook.

6.3 Missed ads by political advertisers

Through our classification from subsection 6.2.1, we can identify advertisers that are known or self-declare to be political actors. For these advertisers, Facebook’s ad policy explicitly mandates that *all* their ads be declared (“ads made by” a political actor). If these advertisers fail to disclose all their ads, Facebook should detect them. With this premise, we can measure whether those pages had any undetected ads that Facebook’s enforcement missed within 14 days after the ad’s first activity. By taking a more holistic approach where we identify groups of pages and ads where enforcement is required, we avoid introducing any interpretation of our own of what should be a ‘political ad’ [102], as well as errors from machine learning models that detect political ads [34, 96]. In-

Table 4: Distribution of observed (political) ads over identified advertiser classes. Classes do not overlap.

Type of page	Pages		All ads		Political ads		Political ad proportion (%)				% pages with		Detected ads		Detected pages	
	#	%	#	%	#	%	overall	Q_{1-3} per page	100% pol. ads	100% pol. ads	#	%	#	%		
Political	102,617	38.6	2,593,727	7.7	2,476,764	59.1	95.5	100.0	100.0	100.0	80.8	6,010	8.3	2,965	21.3	
Government	9,686	3.6	781,102	2.3	190,057	4.5	24.3	0.0	20.0	92.9	22.7	4,577	6.3	956	6.9	
Issue	20,250	7.6	1,722,973	5.1	585,326	14.0	34.0	0.0	25.0	93.9	22.3	11,245	15.5	1,870	13.5	
Media	26,203	9.9	11,814,593	34.9	318,939	7.6	2.7	0.0	0.0	50.0	12.2	24,395	33.6	2,858	20.6	
Other with authorization	30,108	11.3	1,954,851	5.8	604,897	14.4	30.9	44.4	100.0	100.0	52.9	11,073	15.2	1,632	11.7	
Other with detected political ads	3,619	1.4	6,540,777	19.3	15,378	0.4	0.2	0.8	3.8	14.3	4.9	15,378	21.2	3,619	26.0	
Other with only observed non-political ads	47,248	17.8	8,420,746	24.9	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	0.0	
Other without observed ads	26,093	9.8	0	0.0	0	0.0	0.0	—	—	—	0.0	0	0.0	0	0.0	
All with authorization	139,861	52.6	8,132,939	24.0	4,151,486	99.0	51.0	89.5	100.0	100.0	69.4	32,803	45.1	6,412	46.1	
All with declared or detected political ads	143,657	54.0	21,511,759	63.6	4,191,361	100.0	19.5	80.0	100.0	100.0	66.3	72,678	100.0	13,900	100.0	
All pages	265,824	100.0	33,828,769	100.0	4,191,361	100.0	12.4	0.0	80.3	100.0	44.3	72,678	100.0	13,900	100.0	

Table 5: Ads from clear political advertisers missed by Facebook’s enforcement (false negatives). Groups may overlap.

Political page list	Country	Undetected ads		Pages with ≥ 1 undetected ad		Detected ads	
		#	%*	#	%	#	%§
Ad Spending Tracker	US	33	0.01	16	1.93	1	0.00
FEC-registered organizations	US	1,035	0.17	42	2.76	40	0.01
OpenSecrets.org committees	US	129	0.03	24	1.74	14	0.00
CNPJ-registered entities	BR	7,607	1.57	2,038	11.99	668	0.14
CrowdTangle Explore lists	Int.	14,263	2.15	649	10.46	306	0.05
Facebook Page categories	Int.	103,808	4.33	16,116	16.53	5,764	0.25
All political pages		116,963	4.51	16,875	16.44	6,010	0.24
<i>U.S.-only</i>	US	10,940	0.96	1,187	4.75	416	0.04

* False negative rate ($FN/(FN + TP)$) § Proportion of ads detected by Facebook across all ads labeled as political by Facebook within the respective group of pages (TP).

stead, by selecting pages that we believe to be clearly in scope of Facebook’s political ad policy, we have greater confidence that we observe genuine errors in Facebook’s enforcement.

Table 5 shows the different lists of political pages that we derived from the external data sources. To ensure the precision of these lists, one author manually verified all pages with undetected ads from the three U.S.-based lists and removed entries that were not political actors. One such example was a media page that ran a few one-off ads on behalf of a presidential campaign and disclosed them using the campaign’s FEC ID. For the larger list of advertisers that self-declare a political Facebook page category, such as ‘Politician’ or ‘Political Organization,’ we randomly sampled 50 pages from the U.S. (to ensure interpretability) for which one author manually confirmed that all 50 pages belonged to core political actors. We conclude that these lists are a reliable source of political pages that are required under Facebook’s policy to declare all of their ads as political.

Across federal and state-level U.S. political advertisers, i.e., those that Facebook included in its Spending Tracker or verified through an FEC ID as well as major committees tracked by OpenSecrets, performance is nearly perfect with at most 0.17% missed ads, depending on the list (Table 5). If we broaden our view to include additional core political actors in the U.S., identified based on manually curated political Explore lists from CrowdTangle (a subsidiary of Facebook) or the category of their Facebook pages, Facebook misses

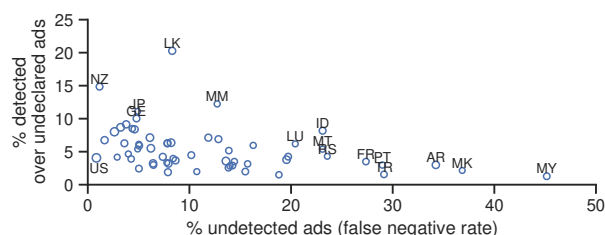


Figure 7: Rates of undetected and detected ads for countries with at least 0.01% of all observed political ads, across pages in political Facebook page categories.

more ads (0.96% in total). As these additional lists also cover advertisers for local elections, the increase in missed ads suggests that Facebook is less successful at identifying smaller political actors. In absolute terms, out of the combined 11,356 ads that U.S. core political actors failed to declare, Facebook was only able to detect for 416 (3.7%) that they were political and therefore enforce its ad policy. Facebook failed to detect the remaining 10,940 electoral ads, including ads from a U.S. senator and former presidential primary candidate with almost 10 million dollar in ad spend, who was able to run undisclosed (and undetected) ads after the U.S. elections.

We now analyze whether Facebook’s enforcement is consistent globally through international lists of political actors, comparing performance in particular to that in the United States. For candidates and parties registered in the 2020 Brazilian municipal elections [98], 1.57% of ads went undetected, even though Facebook could easily match their provided CNPJ ID with the official list of political advertisers. Silva et al. [96] observed a slightly higher false negative rate of 2.2% in the 2018 Brazilian elections, albeit across all advertiser types. Across global lists of political actors curated by CrowdTangle, 2.15% of ads were missed. Across pages that classify themselves in a political Facebook page category, 4.3% of ads were neither declared nor detected. (Note that this is a lower bound, as pages could select generic categories that we do not include here. Within the other lists of political pages, we find 3,311 ‘Public Figure’ and 715 ‘Personal Blog’ pages, for example.) All of these false negative rates are worse

than even the most broadly defined set of political actors in the United States. Across all political advertisers worldwide, we find a false negative rate of 4.5%, almost five times that of the U.S. 16.4% of pages published at least one political ad that was not detected, more than three times as many as in the U.S., even though they are obvious political actors.

We further break down rates of undetected ads for advertisers in political Facebook page categories by country (Figure 7), as these advertisers cover all countries in our scope. Again, Facebook misses the fewest ads in the United States (0.85% false negatives), whereas enforcement can be considerably worse in other countries: in absolute terms, Argentina, Brazil, and India have over 10,000 undetected ads each, while in relative terms, Argentina, North Macedonia, and Malaysia have between 30% and 45% undetected ads. This further suggests that Facebook does not succeed in enforcing its policies consistently worldwide, leaving some users more exposed to violating political ads. This inconsistency may be due to language-specific model deficiencies [92], in particular if little training data is available [82]. Additionally, the review model may fail to incorporate different cultural contexts, with certain topics being considered politically sensitive only in some countries, as Facebook itself recognizes [84]. However, other confounding factors exist, such as the importance of individual pages, the reach of ads, or heightened attention due to ongoing elections [61]. These prevent us from reliably attributing performance differences to one or more causes.

Overall, of all 122,973 ads that political actors do not declare, Facebook only succeeds in detecting a very minor share of 6,010 ads or 4.9%, while failing to detect 116,963 ads, meaning that Facebook's enforcement is ineffective at discovering violating ads even from advertisers with a clear political intent. Moreover, if an advertiser properly declares their ad, Facebook does not have to make an enforcement decision, increasing the significance of any error made for an undeclared ad. In terms of potential exposure, the activity period for these undetected ads is similar to that for detected ads (subsection 5.1). While Facebook does not disclose spend and impressions metadata for ads not known to be political, we use the average ad spend and impressions across each page's political ads to estimate that advertisers likely spent between 4.6 and 9.2 million U.S. dollars on undetected ads, and that these ads had between 2.2 and 2.4 billion impressions. As with detected ads, users are therefore exposed in great quantities to these violating ads. However, undetected ads lack even the transparency that comes with after-the-fact detection, since these ads disappear from the Ad Library and can no longer be scrutinized once they become inactive.

Summary Detection led 22 of the 75 pages with the most detected ads to correctly declare more political ads. However, the remaining pages can and do continue publishing undeclared political ads even after Facebook detected their undeclared ads, and after Facebook prohibited political ads in the U.S. Facebook's policies appear to result in non-political

pages being listed in the Ad Library Report as having political ads, which may have been inadvertently published or erroneously detected. Conversely, we consider only 39% of pages to be core political actors, who we expect to declare all their ads or else get detected by Facebook. Unfortunately, we find at least 116,963 ads from these clearly political advertisers that were missed by Facebook's detection (*false negatives*). Moreover, these missed ads are unevenly distributed worldwide: while for U.S. advertisers only 0.85% of ads are missed, we see a false negative rate of up to 45% in other countries. Put differently, for political pages only 4.9% of undeclared ads are detected, resulting in at least 2.2 billion (estimated) impressions that expose users to ads that hide their political nature and avoid disclosing who paid for the ad.

7 Discussion

Across the ads where Facebook has to make an enforcement decision, we observe 61% more undetected ads (across political pages; subsection 6.3) than detected ads (across all pages; subsection 5.1) within 14 days after their first activity. In addition, we observe that 55% of detections in the U.S. are likely false (subsection 5.2). Translated into classification metrics, we estimate a precision of 0.45, a recall of 0.22, and subsequently an F_1 score of 0.29, all indicative of insufficiently accurate classification, and calling into question whether Facebook's enforcement is truly effective.

Incidentally, these error estimates are conservative and biased favorably towards better performance by Facebook. We quantify false negatives only across clearly political pages, where all ads must be declared, and enforcement is likely easier. If Facebook were to implement detection of every ad from such pages, our conservative estimate of the false negative rate would become zero. However, our estimate does not include potentially missed ads from other (non-political) advertisers; if these publish political ads without disclosing them, the ads would likely be more difficult to detect, given that they must be evaluated individually. Sosnovik and Goga [102] found 4% of 'strong political' ads to be undeclared, similar to our false negative rate. However, 7% of 'political' and 26% of 'marginally political' ads were also undeclared, and such ads were more often placed by NGOs, advocacy groups and charity organizations. This suggests that we would also find a non-zero false negative rate if we were to extend our estimate to include individual ads from non-political actors. Nevertheless, our results present a baseline for ads currently missed by Facebook. Conversely, we conservatively estimate worldwide false positives based on our findings among advertisers in the United States. We found that the false negative rate was lowest in the United States; if a similar trend holds for false positives, Facebook's worldwide false positive rate is likely higher than in the U.S. Therefore, Facebook's global performance is likely worse than our estimate.

We now discuss five factors that enable more effective

enforcement, and highlight how our findings suggest that Facebook’s implementation is lacking in these areas. We also outline recommendations to Facebook for improving its enforcement and reduce erroneously missed or detected ads, as well as improve researchers’ ability to audit its enforcement.

First, in terms of technical capability, Facebook’s enforcement **approach** appears insufficient for the task of classifying political ads. Its automated moderation systems apparently do not learn or incorporate obvious signals of political intent, such as a page’s self-categorization (subsection 6.3), and this despite a high false positive rate (subsection 5.2). Even if Facebook’s review were more performant, the scale of its ad business means that low error rates still result in large absolute counts of missed political ads. *Recommendation:* Facebook should expand its enforcement approach to take the advertiser into account, e.g., by monitoring pages in political categories more strictly [36, 64]. Such simple, clearly enforceable rules could complement the current automated review.

Second, policy enforcement should be timely and come with appropriate **consequences** to prevent future violations. However, pages were still able to repeatedly run undeclared ads, even during the pause on political ads in the U.S. (subsection 6.1). Moreover, violating ads sometimes run for a long period or are already inactive by the time of detection, resulting in large exposure (billions of impressions) before they are caught, if ever (subsection 5.1). *Recommendation:* Facebook should ensure stricter consequences for repeatedly violating advertisers, such as (temporarily) restricting them from running ads.

Third, enforcement must be **consistent** in order to be fair and effective for all users and advertisers. However, next to the overall enforcement errors that suggest inadequate reviewing resources, it appears that missed ads are more common outside the United States, where Facebook’s enforcement suffers from higher false negative rates (subsection 6.3), leaving users there more vulnerable to obscured political ads. *Recommendation:* Facebook should ensure consistent performance globally, independent of an ad’s language. To capture cultural differences, they should engage with local governments, regulators and organizations to adapt policies and enforcement strategies to the local context [64, 101]. This includes identifying country-specific sensitive topics. Furthermore, they should mandate ad declaration worldwide, to ensure that no users in any country are unnecessarily left vulnerable to malicious political advertisers [84, 101].

Fourth, enforcement errors could result from insufficient **ad policies**. We find many largely non-political advertisers who appear to (possibly unintentionally) violate these policies and have detected ads (subsubsection 6.2.2), even though they might have good intentions and be unaware that their ad was ‘political.’ This may be due to ambiguity in ad policies, in particular whether social issue ads “seek to influence public opinion.” Our ‘expert’ annotators did not always agree on whether an ad was political (subsection 5.2), suggesting that

advertisers may also find this difficult, in particular as policies are spread out across many resources [1, 4, 5, 7, 11–13, 18, 27, 28, 48, 56, 57]. *Recommendation:* Facebook should further clarify and simplify its political ad policies, making it very obvious whether an ad is in scope or not [22, 101]. In addition, policies should be collected in one easily discoverable location [67, 78], with updates being clearly indicated and previous versions remaining available [22].

Finally, the quality of enforcement also affects the **transparency** into the political ad ecosystem that the Ad Library is meant to provide. Missed political ads disappear from the Ad Library once they become inactive, and additional metadata such as its spend and impressions are unavailable. Conversely, falsely detected non-political ads result in unrelated advertisers and ads appearing in the Ad Library, which may result in overestimating political ads on Facebook, and increases the (infrastructural and human) resources required to retrieve, process and analyze data for all advertisers (subsection 4.5). *Recommendation:* Although we commend Facebook for their current transparency efforts, as they enable our audit and allow us to suggest improvements, they should expand transparency by including all ads in their archive and API to enable reproducible and scalable analysis of their enforcement [2, 26, 35, 39, 59, 62, 64, 65, 71, 88, 110]. For detected ads, they should also disclose which policy was violated and how they determined this [22], instead of the current binary signal.

However, changes to enforcement and transparency should be balanced with legitimate commercial and privacy concerns around sharing ad metadata, as well as consider adversarial counteractions from advertisers, who for example could attempt to evade efforts to identify them as a political actor (e.g., by selecting an unrelated page category). Actors beyond Facebook may therefore also need to intervene: legislators could harmonize definitions of both political and issue ads across platforms [67] as well as set enforcement and transparency requirements [2, 26] that would be overseen by regulators [110] (Appendix D). Ultimately, such regulatory pressure would entail a shift away from the current self-regulatory model to co-regulation [39, 110]: being allowed to self-regulate policies requires being able to enforce them well, which we show Facebook currently fails to achieve.

8 Conclusion

Through a large-scale collection of all ads from 215,030 pages with political ads over seven months, we conduct an audit of Facebook’s political ad policy enforcement. We study whether this enforcement prevents negligent or malicious advertisers from weakening the integrity of the online political ad ecosystem by running political ads without disclosing them as required, and whether enforcement unnecessarily harms legitimate advertisers. Unfortunately, we find that Facebook’s detection of political ads is flawed: Facebook misses more ads than they detect, and over half of those detected ads are

incorrectly flagged. This enables advertisers to violate policies for an extended time or even evade bans on political ads. We attribute these flaws to limitations in Facebook’s approach that does not sufficiently take into account the political intent of advertisers, allows pages to continue running violating ads, does not appear to be localized well in many countries, and is based on ambiguous policies that are harder to comply with and to enforce. These flaws then result in worse transparency into the online political ad ecosystem, as undisclosed and undetected political ads are neither accounted for in the summary statistics of the Ad Library Report, nor archived in the Ad Library so that the ads could be scrutinized after they become inactive. Yet, despite its flaws, it is also due to this transparency that we can audit Facebook’s enforcement and formulate our recommendations to improve it: By being able to hold platforms accountable, we can work towards more secure online political speech.

Acknowledgments

We thank the anonymous reviewers and our shepherd Mainack Mondal for their valuable feedback. We also thank Davy Preuveneers for help with the classification metrics. Victor Le Pochat holds a PhD Fellowship of the Research Foundation Flanders - FWO (11A3421N). This research is partially funded by the Research Fund KU Leuven, and by the Flemish Research Programme Cybersecurity. Cybersecurity for Democracy at NYU’s Center for Cybersecurity has been supported by Democracy Fund, Luminare, Media Democracy Fund, the National Science Foundation under grant 1814816, Reset, and Wellspring. This material is based upon work supported by the Google Cloud Research Credits program.

References

- [1] *5 Things to Remember About Political and Issue Advertising Around the US 2020 Election*. Facebook for Business. Oct. 26, 2020. URL: <https://www.facebook.com/business/news/facebook-ads-restriction-2020-us-election>.
- [2] *A comprehensive plan to innovate democracy in Europe: Civil society vision for the European Democracy Action Plan*. European Partnership for Democracy, Sept. 2020. URL: <https://epd.eu/wp-content/uploads/2020/09/a-civil-society-vision-for-the-european-democracy-action-plan-input-paper.pdf>.
- [3] *A Look at Facebook and US 2020 Elections*. Report. Facebook, Dec. 2020. URL: <https://about.fb.com/wp-content/uploads/2020/12/US-2020-Elections-Report.pdf>.
- [4] *About Ads About Social Issues, Elections or Politics*. Facebook Business Help Center. Mar. 3, 2021. URL: <https://www.facebook.com/business/help/167836590566506>.
- [5] *About Social Issues*. Facebook Business Help Center. Apr. 7, 2021. URL: <https://www.facebook.com/business/help/214754279118974>.
- [6] *About the Ad Library*. Facebook Business Help Center. Sept. 2, 2020. URL: <https://www.facebook.com/business/help/2405092116183307>.
- [7] *Ad Authorization Exemptions and How They Work*. Facebook Business Help Center. Dec. 15, 2020. URL: <https://www.facebook.com/business/help/387111852028957>.
- [8] *Ad Library*. Facebook. 2021. URL: <https://www.facebook.com/ads/library/>.
- [9] *Ad Library API*. Facebook. 2021. URL: <https://www.facebook.com/ads/library/api/>.
- [10] *Ad Library Report*. Facebook. 2021. URL: <https://www.facebook.com/ads/library/report/>.
- [11] *Ads About Social Issues, Elections or Politics*. Facebook Business Help Center. Mar. 10, 2021. URL: <https://www.facebook.com/business/help/1838453822893854>.
- [12] *Advertising Policies: Restricted Content: Ads About Social Issues, Elections or Politics*. Facebook. 2021. URL: https://www.facebook.com/policies/ads/restricted_content/political.
- [13] *Advertising Policies: Restricted Content: Disclaimers for Ads About Social Issues, Elections or Politics*. Facebook. 2021. URL: https://www.facebook.com/policies/ads/restricted_content/disclaimers.
- [14] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. “Discrimination through Optimization: How Facebook’s Ad Delivery Can Lead to Biased Outcomes”. In: *Proceedings of the ACM on Human-Computer Interaction* 3.CSCW (Nov. 2019). DOI: [10.1145/3359301](https://doi.org/10.1145/3359301).
- [15] Muhammad Ali, Piotr Sapiezynski, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. “Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging”. In: *14th ACM International Conference on Web Search and Data Mining*. WSDM ’21. 2021, pp. 13–21. DOI: [10.1145/3437963.3441801](https://doi.org/10.1145/3437963.3441801).
- [16] Athanasios Andreou, Marcio Silva, Fabrício Benevenuto, Oana Goga, Patrick Loiseau, and Alan Mislove. “Measuring the Facebook Advertising Ecosystem”. In: *2019 Network and Distributed System Security Symposium*. NDSS ’19. 2019. DOI: [10.14722/ndss.2019.23280](https://doi.org/10.14722/ndss.2019.23280).
- [17] *Automated Data Collection Terms*. Facebook. Apr. 15, 2010. URL: https://www.facebook.com/apps/site_scraping_tos_terms.php.
- [18] *Availability for Ads About Social Issues, Elections or Politics*. Facebook Business Help Center. Mar. 11, 2021. URL: <https://www.facebook.com/business/help/2150157295276323>.
- [19] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. “The Menlo Report”. In: *IEEE Security & Privacy Magazine* 10.2 (Mar. 2012), pp. 71–75. DOI: [10.1109/msp.2012.52](https://doi.org/10.1109/msp.2012.52).
- [20] *Beginner’s Guide to Using NCCS Data*. Urban Institute, National Center for Charitable Statistics. Dec. 2018. URL: <https://nccs.urban.org/sites/default/files/2018-12/Guide%20to%20Using%20NCCS%20Data.pdf>.
- [21] Davide Beraldo, Stefania Milan, Jeroen de Vos, Claudio Agosti, Bruno Nadalic Sotic, Rens Vliegthart, Sanne Kruike-meier, Lukas P Otto, Susan A. M. Vermeer, Xiatong Chu, and Fabio Votta. “Political advertising exposed: tracking Facebook ads in the 2021 Dutch elections”. In: *Internet Policy Review* (Mar. 11, 2021). URL: <https://policyreview.info/articles/news/political-advertising-exposed-tracking-facebook-ads-2021-dutch-elections/1543>.
- [22] Sarah E. Bolden, Brian McKernan, and Jennifer Stromer-Galley. *Facebook Political Advertising Transparency Report*. Illuminating. Oct. 5, 2020. URL: <https://news.illuminating.ischool.syr.edu/2020/10/06/facebook-political-advertising-transparency-report/>.

- [23] Grazia Cecere, Clara Jean, Vincent Lefrere, and Catherine E. Tucker. *Tradeoffs in Automated Political Advertising Regulation: Evidence from the COVID-19 Pandemic*. 2020. DOI: 10.2139/ssrn.3603341. SSRN: 3603341.
- [24] *Choosing a Special Ad Category*. Facebook Business Help Center. Feb. 19, 2021. URL: <https://www.facebook.com/business/help/298000447747885>.
- [25] Jason Chuang. *Data Collection Log – EU Ad Transparency Report*. Mozilla. 2019. URL: <https://adtransparency.mozilla.org/eu/log/>.
- [26] *Civil society coalition calls for meaningful transparency on all ads*. European Partnership for Democracy, Sept. 8, 2020. URL: <https://epd.eu/2020/09/08/pressreleaseadstransparency/>.
- [27] *Comply With Local Laws Governing Ads About Social Issues, Elections or Politics*. Facebook Business Help Center. Apr. 10, 2020. URL: <https://www.facebook.com/business/help/534095800360404>.
- [28] *Create Disclaimers and Link Ad Accounts*. Facebook Business Help Center. Jan. 5, 2021. URL: <https://www.facebook.com/business/help/488070228549681>.
- [29] CrowdTangle Team. *CrowdTangle*. 2021. URL: <https://www.crowdtangle.com/>.
- [30] Cybersecurity for Democracy. *Ad Observer*. 2021. URL: <https://adobserver.org/>.
- [31] Tom Dobber, Sanne Kruikemeier, Ellen P. Goodman, Natali Helberger, and Sophie Minihold. *Effectiveness of Online Political Ad Disclosure Labels: Empirical Findings*. Institute for Information Law, University of Amsterdam; Rutgers Institute for Information Policy and Law, Mar. 8, 2021. URL: https://www.uva-icds.net/wp-content/uploads/2021/03/Summary-transparency-disclosure-experiment_update.pdf.
- [32] Paul Duke. *What it's like to actually use Facebook's ad transparency tools*. Online Political Transparency Project. Oct. 30, 2020. URL: <https://medium.com/online-political-transparency-project/7accf22f4ba7>.
- [33] Laura Edelson. *Audit of Facebook ad transparency finds missed political ads*. Online Political Transparency Project. Oct. 22, 2020. URL: <https://medium.com/online-political-transparency-project/603f95027cc6>.
- [34] Laura Edelson. *Publishing Facebook ad data (redux)*. Online Political Transparency Project. Nov. 13, 2020. URL: <https://medium.com/online-political-transparency-project/ff071c41c12e>.
- [35] Laura Edelson, Erika Franklin Fowler, and Jason Chuang. "We need universal digital ad transparency now". In: *TechCrunch* (Oct. 16, 2020). URL: <https://techcrunch.com/2020/10/16/we-need-universal-digital-ad-transparency-now/>.
- [36] Laura Edelson, Tobias Lauinger, and Damon McCoy. "A Security Analysis of the Facebook Ad Library". In: *2020 IEEE Symposium on Security and Privacy*. SP '20. 2020, pp. 661–678. DOI: 10.1109/SP40000.2020.00084.
- [37] Laura Edelson, Shikhar Sakhuja, Ratan Dey, and Damon McCoy. *An Analysis of United States Online Political Advertising Transparency*. 2019. arXiv: 1902.04385 [cs.SI].
- [38] *Election Integrity at Facebook*. Facebook for Business. 2021. URL: <https://www.facebook.com/business/m/election-integrity>.
- [39] *ERGA Report on disinformation: Assessment of the implementation of the Code of Practice*. European Regulators Group for Audiovisual Media Services, May 2020. URL: <https://erga-online.eu/wp-content/uploads/2020/05/ERGA-2019-report-published-2020-LQ.pdf>.
- [40] *Exempt Organizations Business Master File Extract*. Internal Revenue Service. 2021. URL: <https://www.irs.gov/charities-non-profits/exempt-organizations-business-master-file-extract-eo-bmf>.
- [41] *Facebook Ads Library Assessment*. Office of the French Ambassador for Digital Affairs. 2020. URL: <https://disinfo.quaidorsay.fr/en/facebook-ads-library-assessment>.
- [42] *Facebook response to the public consultation for the European Democracy Action Plan*. Facebook, Sept. 15, 2020. URL: https://about.fb.com/wp-content/uploads/sites/10/2020/09/Facebook_Response_European_Democracy_Action_Plan_2020.09.15.pdf.
- [43] Irfan Faizullahoy and Aleksandra Korolova. "Facebook's Advertising Platform: New Attack Vectors and the Need for Interventions". In: *2nd Workshop on Technology and Consumer Protection*. ConPro '18. 2018. URL: <https://www.ieee-security.org/TC/SPW2018/ConPro/papers/faizullahoy-conpro18.pdf>.
- [44] Erika Franklin Fowler, Michael M. Franz, and Travis N. Ridout. "Online Political Advertising in the United States". In: *Social Media and Democracy: The State of the Field, Prospects for Reform*. Ed. by Nathaniel Persily and Joshua A. Tucker. SSRN Anxieties of Democracy. Cambridge University Press, 2020, pp. 111–138. ISBN: 9781108890960. DOI: 10.1017/9781108890960.
- [45] Sarah Frier. "Facebook's Political Rule Blocks Ads for Bush's Beans, Singers Named Clinton". In: *Bloomberg* (July 2, 2018). URL: <https://www.bloomberg.com/news/articles/2018-07-02/facebook-s-algorithm-blocks-ads-for-bush-s-beans-singers-named-clinton>.
- [46] Devin Gaffney and J. Nathan Matias. "Caveat emptor, computational social science: Large-scale missing data in a widely-published Reddit corpus". In: *PLOS ONE* 13.7 (July 2018). DOI: 10.1371/journal.pone.0200162.
- [47] David Gale. "Facebook's Problem With Veterans". In: *The Wall Street Journal* (Aug. 7, 2018). URL: <https://www.wsj.com/articles/facebook-s-problem-with-veterans-1533682511>.
- [48] *Get Authorized to Run Ads About Social Issues, Elections or Politics*. Facebook Business Help Center. Mar. 3, 2021. URL: <https://www.facebook.com/business/help/208949576550051>.
- [49] Avijit Ghosh, Giridhari Venkatadri, and Alan Mislove. "Analyzing Political Advertisers' Use of Facebook's Targeting Features". In: *3rd Workshop on Technology and Consumer Protection*. ConPro '19. 2019. URL: <https://www.ieee-security.org/TC/SPW2019/ConPro/papers/ghosh-conpro19.pdf>.
- [50] José González Cabañas, Ángel Cuevas, Aritz Arrate, and Rubén Cuevas. "Does Facebook Use Sensitive Data for Advertising Purposes?" In: *Communications of the ACM* 64.1 (Dec. 2020), pp. 62–69. DOI: 10.1145/3426361.
- [51] José González Cabañas, Ángel Cuevas, and Rubén Cuevas. "FDVT: Data Valuation Tool for Facebook Users". In: *2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. 2017, pp. 3799–3809. DOI: 10.1145/3025453.3025903.
- [52] José González Cabañas, Ángel Cuevas, and Rubén Cuevas. "Unveiling and Quantifying Facebook Exploitation of Sensitive Personal Data for Advertising Purposes". In: *27th USENIX Security Symposium*. USENIX Security '18. 2018, pp. 479–495.
- [53] Jessica Guynn. "Facebook takes down ads mentioning African-Americans and Hispanics, calling them political". In: *USA TODAY* (Oct. 17, 2018). URL: <https://www.usatoday.com/story/news/2018/10/17/facebook-labels-african-american-hispanic-mexican-ads-political/1608841002/>.
- [54] Jeroen van der Ham and Roland van Rijswijk-Deij. "Ethics and Internet Measurements". In: *Journal of Cyber Security and Mobility* 5.4 (2017), pp. 287–308. DOI: 10.13052/jcsm2245-1439.543.

- [55] Joris van Hoboken, Naomi Appelman, Ronan Ó Fathaigh, Paddy Leerssen, Tarlach McGonagle, Nico van Eijk, and Natali Helberger. *The legal framework on the dissemination of disinformation through Internet services and the regulation of political advertising*. Institute for Information Law, University of Amsterdam, Dec. 2019. URL: https://www.ivir.nl/publicaties/download/Report_Disinformation_Dec2019-1.pdf.
- [56] *How Ads About Social Issues, Elections or Politics Are Reviewed (With Examples)*. Facebook Business Help Center. Mar. 3, 2021. URL: <https://www.facebook.com/business/help/313752069181919>.
- [57] *How Disclaimers Work for Ads About Social Issues, Elections or Politics*. Facebook Business Help Center. Aug. 19, 2020. URL: <https://www.facebook.com/business/help/198009284345835>.
- [58] *How do I choose to see fewer ads about social issues, elections or politics on Facebook?* Facebook Help Center. 2021. URL: <https://www.facebook.com/help/595432167810439>.
- [59] Philip N. Howard. “A Way to Detect the Next Russian Misinformation Campaign”. In: *The New York Times* (Mar. 27, 2019). URL: <https://www.nytimes.com/2019/03/27/opinion/russia-elections-facebook.html>.
- [60] Basileal Imana, Aleksandra Korolova, and John Heidemann. “Auditing for Discrimination in Algorithms Delivering Job Ads”. In: *The Web Conference 2021*. WWW ’21. 2021, pp. 3767–3778. DOI: 10.1145/3442381.3450077.
- [61] *Inquiry into and report on all aspects of the conduct of the 2019 Federal Election and matters related thereto*. Submission to the Joint Standing Committee on Electoral Matters 140.1. Facebook, Sept. 16, 2020. URL: <https://www.aph.gov.au/DocumentStore.ashx?id=acd7119e-7cc7-474f-ac60-6af6c5c3d933&subId=671217>.
- [62] Karolina Iwańska, Katarzyna Szymielewicz, Dominik Batorski, Maciej Baranowski, Jakub Krawiec, Krzysztof Izdebski, and Daniel Macyszyn. *Who (really) targets you? Facebook in Polish election campaigns*. Panoptikon Foundation, ePaństwo Foundation, SmartNet Research & Solutions, Mar. 17, 2020. URL: <https://panoptikon.org/political-ads-report>.
- [63] Amelia M. Jamison, David A. Broniatowski, Mark Dredze, Zach Wood-Doughty, DureAden Khan, and Sandra Crouse Quinn. “Vaccine-related advertising in the Facebook Ad Archive”. In: *Vaccine* 38.3 (2020), pp. 512–520. DOI: 10.1016/j.vaccine.2019.10.066.
- [64] Julian Jaurisch. *Defining Online Political Advertising. How Difficulties in Delineating Paid Political Communication Can Be Addressed*. Stiftung Neue Verantwortung, Nov. 2020. URL: https://www.stiftung-nv.de/sites/default/files/snv_definingpoliticalads.pdf.
- [65] Sam Jeffers. *How to take a “gold standard” approach to political advertising transparency and policy*. Who Targets Me. Oct. 28, 2020. URL: <https://whotargets.me/nl/how-to-take-a-gold-standard-approach-to-political-advertising-transparency-and-policy/>.
- [66] Simon Kemp. *Digital 2021: Global Overview Report*. Hootsuite & We Are Social, Jan. 27, 2021. URL: <https://datareportal.com/reports/digital-2021-global-overview-report>.
- [67] Daniel Kreiss and Bridget Barrett. “Democratic Tradeoffs: Platforms and Political Advertising”. In: *Ohio State Technology Law Journal* 16.2 (2020), pp. 493–519.
- [68] Klaus Krippendorff. *Content analysis: an introduction to its methodology*. 4th ed. Sage, 2018. ISBN: 9781506395661.
- [69] Anja Lambrecht and Catherine Tucker. “Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads”. In: *Management Science* 65.7 (July 2019), pp. 2966–2981. DOI: 10.1287/mnsc.2018.3093.
- [70] Issie Lapowsky. “How right-wing websites are getting around Facebook’s political ad ban”. In: *Protocol* (Dec. 11, 2020). URL: <https://www.protocol.com/facebook-political-ad-ban-news>.
- [71] Paddy Leerssen, Jef Ausloos, Brahim Zarouali, Natali Helberger, and Claes H. de Vreese. “Platform ad archives: promises and pitfalls”. In: *Internet Policy Review* 8.4 (Oct. 2019). DOI: 10.14763/2019.4.1421.
- [72] Rachel Lerman and Cat Zakrzewski. “Facebook ban on new political ads starts off with major hiccups”. In: *The Washington Post* (Oct. 27, 2020). URL: <https://www.washingtonpost.com/technology/2020/10/27/facebook-ban-new-political-ads/>.
- [73] Kevin Macnish and Jeroen van der Ham. “Ethics in cybersecurity research and practice”. In: *Technology in Society* 63, 101382 (Nov. 2020). DOI: 10.1016/j.techsoc.2020.101382.
- [74] Aaron Mak. “Facebook Thought an Ad From Bush’s Baked Beans Was “Political” and Removed It”. In: *Slate* (May 30, 2018). URL: <https://slate.com/technology/2018/05/bushs-baked-beans-fell-victim-to-facebooks-political-ads-system.html>.
- [75] J. Nathan Matias, Austin Hounsel, and Nick Feamster. *Software-Supported Audits of Decision-Making Systems: Testing Google and Facebook’s Political Advertising Policies*. 2021. arXiv: 2103.00064 [cs.HC].
- [76] J. Nathan Matias, Austin Hounsel, and Melissa Hopkins. “We Tested Facebook’s Ad Screeners and Some Were Too Strict”. In: *The Atlantic* (Nov. 2, 2018). URL: <https://www.theatlantic.com/technology/archive/2018/11/do-big-social-media-platforms-have-effective-ad-policies/574609/>.
- [77] *Media Bias/Fact Check*. Media Bias Fact Check, LLC. 2021. URL: <https://mediabiasfactcheck.com/>.
- [78] Carly Miller. “Facebook, It’s Time to Put the Rules in One Place”. In: *Lawfare* (Mar. 5, 2021). URL: <https://www.lawfareblog.com/facebook-its-time-put-rules-one-place>.
- [79] Blake Montgomery. “Facebook Axed Pro-Vaccine Ads, Let Anti-Vaxxer Conspiracies Slip Through”. In: *The Daily Beast* (Oct. 25, 2019). URL: <https://www.thedailybeast.com/facebook-axed-pro-vaccine-ads-from-hospitals-and-health-orgs-let-anti-vaxxer-ads-slip-through>.
- [80] Blake Montgomery. “Facebook’s Moderators Took Down the Tech Giant’s Own Pro-Equality Ads”. In: *The Daily Beast* (Mar. 10, 2021). URL: <https://www.thedailybeast.com/facebook-moderation-system-took-down-the-tech-giants-own-black-history-month-ads>.
- [81] *National Center for Charitable Statistics Data Archive*. Urban Institute, National Center for Charitable Statistics. 2021. URL: <https://nccs-data.urban.org/index.php>.
- [82] Wilhelmina Nekoto et al. “Participatory Research for Low-resourced Machine Translation: A Case Study in African Languages”. In: *Findings of the Association for Computational Linguistics*. EMNLP ’20. Nov. 2020, pp. 2144–2160. DOI: 10.18653/v1/2020.findings-emnlp.195.
- [83] *NewsGuard*. NewsGuard Technologies. 2021. URL: <https://www.newsguardtech.com/>.
- [84] *Online Political Ads: A study of inequality in transparency standards*. Privacy International, Jan. 2021. URL: https://privacyinternational.org/sites/default/files/2021-01/AdsTransparency_TOPUBLISH.pdf.

- [85] *OpenSecrets*. OpenSecrets. 2021. URL: <https://www.opensecrets.org/>.
- [86] Asha Phillips. *Explore Lists in CrowdTangle*. 2021. URL: <https://help.crowdtangle.com/en/articles/2450455>.
- [87] Filipe N. Ribeiro, Koustuv Saha, Mahmoudreza Babaei, Lucas Henrique, Johnnatan Messias, Fabrício Benevenuto, Oana Goga, Krishna P. Gummadi, and Elissa M. Redmiles. “On Microtargeting Socially Divisive Ads: A Case Study of Russia-Linked Ad Campaigns on Facebook”. In: *2nd Conference on Fairness, Accountability, and Transparency*. FAT* ’19. 2019, pp. 140–149. DOI: 10.1145/3287560.3287580.
- [88] Aaron Rieke and Miranda Bogen. *Leveling the Platform: Real Transparency for Paid Messages on Facebook*. Upturn, May 2018. URL: <https://www.upturn.org/reports/2018/facebook-ads/>.
- [89] Guy Rosen, Katie Harbath, Nathaniel Gleicher, and Rob Leathern. *Helping to Protect the 2020 US Elections*. Facebook. Oct. 21, 2019. URL: <https://about.fb.com/news/2019/10/update-on-election-integrity-efforts/>.
- [90] Eli Rosenberg. “Facebook blocked many gay-themed ads as part of its new advertising policy, angering LGBT groups”. In: *The Washington Post* (Oct. 3, 2018). URL: <https://www.washingtonpost.com/technology/2018/10/03/facebook-blocked-many-gay-themed-ads-part-its-new-advertising-policy-angering-lgbt-groups/>.
- [91] Matthew Rosenberg. “Ad Tool Facebook Built to Fight Disinformation Doesn’t Work as Advertised”. In: *The New York Times* (July 25, 2019). URL: <https://www.nytimes.com/2019/07/25/technology/facebook-ad-library.html>.
- [92] Justin Scheck, Newley Purnell, and Jeff Horwitz. “Facebook Employees Flag Drug Cartels and Human Traffickers. The Company’s Response Is Weak, Documents Show.” In: *The Wall Street Journal* (Sept. 16, 2021). URL: <https://www.wsj.com/articles/facebook-drug-cartels-human-traffickers-response-is-weak-documents-11631812953>.
- [93] Mark Scott. “Political ads on Facebook disappear ahead of UK election”. In: *Politico* (Dec. 10, 2019). URL: <https://www.politico.com/news/2019/12/10/political-ads-on-facebook-disappear-ahead-of-uk-election-081376>.
- [94] Mark Scott and Zach Montellaro. “Scores of political groups sidestepped Facebook’s ad ban”. In: *Politico* (Mar. 4, 2021). URL: <https://www.politico.com/news/2021/03/04/political-groups-facebook-ad-ban-473698>.
- [95] Márcio Silva and Fabrício Benevenuto. “COVID-19 Ads as Political Weapon”. In: *36th ACM/SIGAPP Symposium on Applied Computing*. SAC ’21. 2021, pp. 1705–1710. DOI: 10.1145/3412841.3442043.
- [96] Márcio Silva, Lucas Santos de Oliveira, Athanasios Andreou, Pedro Olmo Vaz de Melo, Oana Goga, and Fabrício Benevenuto. “Facebook Ads Monitor: An Independent Auditing System for Political Ads on Facebook”. In: *The Web Conference 2020*. WWW ’20. 2020, pp. 224–234. DOI: 10.1145/3366423.3380109.
- [97] Craig Silverman and Ryan Mac. “Facebook Promised To Label Political Ads, But Ads For Biden, The Daily Wire, And Interest Groups Are Slipping Through”. In: *BuzzFeed News* (Oct. 22, 2020). URL: <https://www.buzzfeednews.com/article/craigsilverman/facebook-biden-election-ads>.
- [98] *Sistema de Consulta a CNPJs*. Tribunal Superior Eleitoral. 2020. URL: <https://spce-cnpj.tse.jus.br/spce2016.cnpj/internet/#/eleicoes>.
- [99] Garrett Sloane. “Facebook repeatedly blocked ads showing wheelchair, says disabilities apparel retailer”. In: *Ad Age* (Dec. 21, 2020). URL: <https://adage.com/article/digital/facebook-repeatedly-blocked-ads-showing-wheelchair-says-disabilities-apparel-retailer/2302536>.
- [100] Sarah Smith. “Facebook denies Houston’s ads promoting fair housing over race, religion references”. In: *Houston Chronicle* (May 15, 2019). URL: <https://www.houstonchronicle.com/news/houston-texas/houston/article/Facebook-denies-Houston-s-ads-promoting-fair-13847765.php>.
- [101] *Social media companies are failing to provide adequate advertising transparency to users globally*. Privacy International, Oct. 3, 2019. URL: https://privacyinternational.org/sites/default/files/2019-10/cop-2019_0.pdf.
- [102] Vera Sosnovik and Oana Goga. “Understanding the Complexity of Detecting Political Ads”. In: *The Web Conference 2021*. WWW ’21. 2021, pp. 2002–2013. DOI: 10.1145/3442381.3450049.
- [103] Michael A. Stelzner. *2020 Social Media Marketing Industry Report*. Social Media Examiners, May 2020. URL: <https://alltomalmhult.se/wp-content/uploads/2020/05/IndustryReport-2020.pdf>.
- [104] Darius Tahir. “Facebook political ad ban blocks pro-vaccine messages”. In: *Politico* (Feb. 21, 2021). URL: <https://www.politico.com/news/2021/02/21/facebook-ad-pro-vaccine-ban-470304>.
- [105] Lori Trahan. *Trahan Leads Introduction of Social Media DATA Transparency Legislation*. May 20, 2021. URL: <https://trahan.house.gov/news/documentsingle.aspx?DocumentID=2112>.
- [106] *US 2020 Elections Ad Spending Tracker*. Facebook. Jan. 23, 2021. URL: <https://about.fb.com/wp-content/uploads/2021/01/US-2020-Elections-Ad-Spending-Tracker.zip>.
- [107] *US Ad Restriction Period*. Facebook. Oct. 2020. URL: https://www.facebook.com/gms_hub/share/facebook-us-ad-restriction-period-overview.pdf.
- [108] Jennifer Valentino-DeVries. “I Approved This Facebook Message – But You Don’t Know That”. In: *ProPublica* (Feb. 13, 2018). URL: <https://www.propublica.org/article/i-approved-this-facebook-message-but-you-dont-know-that>.
- [109] Giridhari Venkatadri, Athanasios Andreou, Yabing Liu, Alan Mislove, Krishna P. Gummadi, Patrick Loiseau, and Oana Goga. “Privacy Risks with Facebook’s PII-Based Targeting: Auditing a Data Broker’s Advertising Interface”. In: *2018 IEEE Symposium on Security and Privacy*. SP ’18. 2018, pp. 89–107. DOI: 10.1109/SP.2018.00014.
- [110] *Virtual Insanity? The need to guarantee transparency in digital political advertising*. European Partnership for Democracy, Mar. 2020. URL: <https://epd.eu/virtual-insanity/>.
- [111] Nancy Watzman. *The political ads Facebook won’t show you*. Cybersecurity for Democracy. May 12, 2021. URL: <https://medium.com/cybersecurity-for-democracy/e0d6181bca25>.
- [112] Madelyn Webb and Bethan John. *We need to know more about political ads. But can transparency be a trap?* First Draft Footnotes. Mar. 25, 2021. URL: <https://medium.com/1st-draft/542df2a52f21>.
- [113] *What is the Facebook Ad Library and how do I search it?* Facebook Help Center. 2021. URL: <https://www.facebook.com/help/259468828226154>.
- [114] André Wheeler. “Facebook under fire after ads for anti-HIV drug PrEP deemed political”. In: *The Guardian* (Oct. 31, 2019). URL: <https://www.theguardian.com/technology/2019/oct/31/facebook-prep-ads-instagram-political>.